

EXPRESS MAIL LABEL NO.: EV 347799604 US DATE OF DEPOSIT: August 6, 2003

I hereby certify that this paper and fee are being deposited with the United States Postal Service Express Mail Post Office to Addressee service under 37 CFR § 1.10 on the date indicated below and is addressed to Mail Stop Patent Applications, Commissioner for Patents, PO Box 1450, Alexandria, VA 22313-1450.

Sandra M. Campbell
NAME OF PERSON MAILING PAPER AND FEE

Sandra M. Campbell
SIGNATURE OF PERSON MAILING PAPER AND FEE

Inventor(s): Allen M. Gilbert
David L. Kaminsky
Balachandar Rajaraman

Online Autonomic Operations Guide

BACKGROUND OF THE INVENTION

Statement of the Technical Field

[0001] The present invention relates to the field of network computing administration and more particularly to network administration using policy based management.

Description of the Related Art

[0002] The task of modern network administration differs significantly from that of days gone by. Not just a decade ago, network administration primarily entailed the addition and deletion of network users, the management of print queues, and the supervision and operation of daily backup procedures. Most if not all resources required by network applications remained present in the network itself, and few if any network applications depended upon the operation of other, co-executing applications. In fact, the notion of an enterprise application, as compared to a mere network application remained largely within the realm of academia as a decade ago, the

enabling technologies had not advanced enough in terms of speed and reliability to facilitate true enterprise computing.

[0003] Much has changed since the early days of network computing. Today, enterprise computing permeates the electronic landscape. While some enterprise applications remain largely stand-alone, most rely in some respect on a co-existing enterprise application or a soft enterprise resource, such as a database application, Web application server, or other cooperating component. Thus, the administration of the system has advanced far beyond user and print queue administration and daily backup routines. Today, the interdependencies among network components presents a significant challenge to the administrator. In this regard, the management of a single network component can depend upon the state of a multiplicity of other network components.

[0004] The task of network management recently has grown to include policy based management principles. Policy based management principles initially included rules for authentication only. Specifically, the rules specified which users defined within the network were permitted to perform which administrative tasks upon which network components. Typically, the authority to perform such administrative tasks represents the sole type of rule managed by policy within the enterprise. Yet, more recent policy based management principles relate more specifically to differentiated service according to the terms of a service level agreement (SLA). Nevertheless, policy based management principles have not addressed other aspects of network administration--particularly those aspects of network administration related to the interoperability of

separate, but interdependent resources, and the ability of the enterprise to behave autonomically, or at least partially autonomically.

[0005] In the famed manifesto, *Autonomic Computing: IBM's Perspective on the State of Information Technology*, Paul Horn, Senior Vice President of IBM Research, observed, "It's not about keeping pace with Moore's Law, but rather dealing with the consequences of its decades-long reign." Given this observation, Horn suggested a computing parallel to the autonomic nervous system of the biological sciences. Namely, whereas the autonomic nervous system of a human being monitors, regulates, repairs and responds to changing conditions without any conscious effort on the part of the human being, in an autonomic computing system, the system must self-regulate, self-repair and respond to changing conditions, without requiring any conscious effort on the part of the computing system operator.

[0006] Thus, while the autonomic nervous system can relieve the human being from the burden of coping with complexity, so too can an autonomic computing system. Rather, the computing system itself can bear the responsibility of coping with its own complexity. The crux of the IBM manifesto relates to eight principal characteristics of an autonomic computing system:

- I. The system must "know itself" and include those system components which also possess a system identify.
- II. The system must be able to configure and reconfigure itself under varying and unpredictable conditions.

- III. The system must never settle for the status quo and the system must always look for ways to optimize its workings.
- IV. The system must be self-healing and capable of recovering from routine and extraordinary events that might cause some of its parts to malfunction.
- V. The system must be an expert in self-protection.
- VI. The system must know its environment and the context surrounding its activity, and act accordingly.
- VII. The system must adhere to open standards.
- VIII. The system must anticipate the optimized resources needed while keeping its complexity hidden from the user.

[0007] Notably, in accordance with the eight tenants of autonomic computing, several single system and peer-to-peer systems have been proposed in which self-configuration, management and healing have provided a foundation for autonomic operation. Yet, despite the eight tenants of autonomic computing, no existing implementation has addressed the need to efficiently and autonomically manage the administration of interdependent components in the enterprise. In particular, what remains at present is a purely manual system governed by hundreds of print pages of operations manuals for managing the network. The manuals generally specify actions to be performed, the timing of those actions, and the persons authorized to perform such actions. Yet, the human enforcement of the rules specified by the operations manual has proven error prone and highly ineffective.

SUMMARY OF THE INVENTION

[0008] The present invention is a systems administration policy enforcement system, method and apparatus. The present invention addresses the deficiencies of conventional systems administration in that the present invention overcomes the human error associated with the multiplicity of rules for administering any one resource in a system. In this regard, rather than requiring an administrator to refer to a voluminous set of printed manuals to determine the requisite state of related resources and the environment in general necessary to perform an administrative task, an administrative policy can be established which embodies the requisite state. A policy enforcement process can permit the administration of a resource in the system only when the requisite state of related resources and the environment comports with the rules of the policy.

[0009] In a method for enforcing the administration policy of a system in accordance with the present invention, a request to perform an administrative task directed to a resource within the computing network can be received. Responsive to receiving the request, an administration policy including a set of rules for governing the administrative task can be retrieved as can state data for the resource. Subsequently, the retrieved policy can be applied to the retrieved state data. Consequently, the administrative task can be permitted only if the retrieved state data satisfies the set of rules in the retrieved policy. Optionally, when a task is not permitted, the initiator of the task can be notified as to why permission to perform the task has been denied.

[0010] Importantly, the administration policy can be established by duly qualified administrators of the system, and in particular, by administrators of the various resources of the system. To that end, a user interface can be provided for establishing the set of rules for the administration policy. Once established, the policy can be stored the administration policy for subsequent retrieval in the retrieving step. The stored policy can consider not only the requisite state of the resource, but also environmental information for the system. The environmental information can range from the identification of the requesting administrator to one or more properties of the system's resources (such as, but not limited to, CPU utilization) to the time of day. In this regard, responsive to the request to administer the resource, the environmental information can be further retrieved for the computing network and the administrative task can be permitted only if the retrieved environmental data satisfies the set of rules in the retrieved policy.

[0011] Notably, the step of further retrieving the state data can include the step of retrieving state data both for the resource and also for other related resources in the computing network. To that end, the administrative task can be disallowed if the further retrieved state data fails to satisfy the set of rules in the retrieved policy either alone or in conjunction with other state data. Once disallowed, a related resource can be identified which has a state which gave rise to the state data for the resource which fails to satisfy the set of rules in the retrieved policy.

[0012] Significantly, upon identifying the related resource, a remediation of the condition in the related resource giving rise to the problematic state can be requested

so that the state of the related resource ultimately satisfies the set of rules in the retrieved policy. Once the condition in the related resource has been remediated, the administrative task can be permitted. Importantly, the steps of disallowing, identifying, requesting and further permitting can be performed autonomically. Moreover, the steps of disallowing, identifying, requesting and further permitting can be performed recursively for each related resource whose state gives rise to a failure of the primary resource to satisfy the retrieved policy.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0013] There are shown in the drawings embodiments which are presently preferred, it being understood, however, that the invention is not limited to the precise arrangements and instrumentalities shown, wherein:
- [0014] Figure 1 is a block illustration of system and method for performing policy based administration of a system in accordance with the inventive arrangements;
- [0015] Figure 2 is a schematic illustration of a policy evaluation component for use in the system of Figure 1; and,
- [0016] Figure 3 is a flow chart illustrating a process for autonomically enforcing an administration policy in accordance with a specific, preferred aspect of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0017] The present invention is a system, method and apparatus for autonomic policy based systems management. In particular, in accordance with the present invention a set of policies can be established for the management of individual components in a system. As defined herein, components can include application components, software resources including application servers and databases relied upon by one or more of the application components, and hardware resources, including physical servers, communications bandwidth, disk storage and the like. Individual policies corresponding to individual components or combinations of individual components can establish rules which can constrain the management of the component.

[0018] The rules defined within the individual policies can include the identities of administrators authorized to undertake particular management operations, the timing of such operations, and most importantly, the requisite state both of the environment and other components required to undertake the particular requested management operation. Where the state of the environment and the other components does not meet the required state specified by the rules in the policy, the management operation can be disallowed. Alternatively, those components whose state inhibits the performance of a requested management operation can be autonomically managed so as to achieve the requisite state.

[0019] Specifically, in the autonomic configuration, when an administrator requests an action which subsequently becomes blocked by a rule in the pertinent

policy, the blocking element can be resolved preferably without intervention by the administrator. As an example, in the event that a rule permits a database shutdown only once the database has undertaken an incremental backup procedure, a request to shutdown the database would become blocked until the backup procedure had occurred. Rather than simply block the request, though, the database can be notified of the requirement to undertake an incremental backup. Once the backup has occurred, the shutdown request can be honored.

[0020] Figure 1 is a block illustration of system and method for performing policy based administration of a system 120 in accordance with the inventive arrangements. The system of the present invention can include a set of rules 150A, 150B, 150n defining an administration policy 140. The policy 140 can specify the identities of administrators authorized to undertake particular management operations, the timing of such operations, and the requisite state both of the environment and other components required to undertake the particular requested management operation. The rules 150A, 150B, 150n in the policy can be established by any administrator and not merely an administrator performing a particular operation at any given time. A policy enforcement processor 300 can enforce the policy 140 responsive to a request to administer elements 130 within the system 120. Notably, such elements can include software and hardware components, operating system components, administrators authorized to manage elements of the enterprise, and the environment generally and of the system 120 itself.

[0021] In operation, an administrator 110 can establish the rules 150A, 150B, 150n of the policy 140. Subsequently, the administrator 110 (or another administrator or other user) can administer the system 120. By administer, it is meant that the administrator 110 can command the operation of one or more of the elements 130 of the system. Such commands can range from startup and shutdown operations, to query and configuration commands. In any case, the policy enforcement processor 300 can receive the administration command issued by the administrator 110. Responsive to the receipt of the administration command, the policy enforcement processor 300 can retrieve the pertinent rules 150A, 150B, 150n of the policy 140 which apply to the requested administration command. Additionally, the policy enforcement processor 300 can query and retrieve the state 160 of the pertinent elements of the system 120. Applying the retrieved pertinent rules 150A, 150B, 150n to the state 160, the policy enforcement processor 300 can determine whether requested administration command ought to be disallowed (blocked), or whether the requested administration command ought to be permitted.

[0022] To facilitate the handling of requests to administer the system 120 of Figure 1, each participating administerable one of the elements 130 (principally software and operating system components) can be configured to route administration requests to a policy enforcement processor, referred to hereafter as a "policy evaluation component". Figure 2 is a schematic illustration of a policy evaluation component 240 which has been configured for use in the system of Figure 1. In accordance with the inventive arrangements, a participating application component 210 can be coupled to an

exit routine 220. The exit routine 220, in particular, can be inserted into the administration console of the application 210. When an administrator initiates an action, an indicator of the action is passed to the policy enforcement component 240 in the form of a request 290. The request 290 can specify not only the nature of the action, but also identifying data suitable for use in authenticating the action in and of itself.

[0023] The policy evaluation component 240 can retrieve all applicable rules 250, for instance using a rules engine (not shown), such engine being well-known in the art. (An example includes the ABLE™ rules engine manufactured by IBM Corporation of Armonk, New York, United States of America) Through the rules engine, the policy evaluation component 240 can review the information required to evaluate the rules 250. The information can include, for instance, information passed on the initial request 290, subsequent queries of the requesting resource and requests to other resources, collectively 270, exogenous factors 280, and history 260, including both history stored at the policy evaluation component 240, and history stored else where in the system. The policy evaluation component 240 can evaluate the rules 250 in combination with the information to determine whether the action is permissible. If the action is permissible, the policy evaluation component 240 can return a result 230 to the exit routine 220 indicating whether the action is permitted.

[0024] Importantly, rather than merely notifying the administrator whether a requested action is permitted, or whether the action has been blocked, the policy evaluation component 240 can request the remediation of the inhibiting elements of the network in an autonomic fashion. In this regard, Figure 3 is a flow chart illustrating a

process for autonomically enforcing a network administration policy in accordance in the policy enforcement processor 300 of Figure 1. Beginning in block 305, a request for administrative action can be received. In block 310, the identification of the administrator can be retrieved from the request and in block 315, the administrator can be authenticated. If the identity of the administrator cannot be authenticated in decision block 320, the requested action can be rejected in block 360 with an associated explanation for the rejection having been concurrently forwarded to the administrator in block 365. Otherwise, the process can continue through block 325.

[0025] In block 325, a set of rules pertaining to the relevant policy can be retrieved for use in evaluating the request. In block 330, a list of state data for the system and other exogenous information, including historical logs and environmental elements (e.g. the time of day) can be constructed. In block 335, the rules can be applied to the list of state data and information to determine whether the action is permissible in view of the current state of the network. In block 340, if it is determined that the action is permissible, in block 345 the action can be permitted and a confirmation can be returned to the administrator in block 350. Otherwise, in block 355 remedial measures can be requested from the element of the network which has given rise to the blockage. Once the blockage has been remediated, the process can continue through blocks 330 through 350. If the blockage cannot be remediated, then in block 360 the action can be rejected and an explanation for the rejection can be forwarded to the administrator in block 365.

[0026] As an example, in the context of a request to shutdown a database, a pertinent set of rules in a network administration policy might specify that the database cannot be shutdown during working hours, while an application server remains connected to the database, prior to the performance of an incremental backup, and by anyone other than a specified database administrator. When an administrator requests a shutdown of the database, the exit routine of the console of the database component can trap the request, forwarding such request to a policy evaluation component. The policy evaluation component can retrieve the set of rules and associated state data and information. The state data and information can include, among other things, the identity of the administrator and the components coupled to the database, the time of day and the state of the database (e.g. the last time an incremental backup had been performed).

[0027] Initially, the policy evaluation component can determine whether the identity of the administrator is that of an authorized database administrator. If not, the request can be rejected and the administrator can be notified that the administrator lacks the relevant credentials to undertake a database shutdown operation. If the time of day falls within the impermissible working hours range, the request again can be rejected and the administrator can be notified that the administrator must wait until after working hours to perform a shutdown. If an application component remains coupled to the database, again the request can be rejected pending the decoupling of the application component from the database. Finally, if the database had not been

incrementally backed up, the action can be rejected and a suitable explanation can be forwarded to the administrator.

[0028] Significantly, in the latter two circumstances, the policy evaluation component can attempt to resolve the blocking condition without significant intervention by the administrator. More particularly, in the case of a coupled application component, the policy evaluation component with or without the explicit approval of the administrator can forward a request to the coupled application to decouple itself from the database. Similarly, in the case of the incremental backup requirement, the policy evaluation component can forward a request to the database to perform an incremental backup. In both cases, the policy evaluation component can suspend the shutdown operation until the request criteria of the policy have been met. At that time, the policy evaluation component can resume the processing of the shutdown request without intervention by the administrator.

[0029] It will be recognized by the skilled artisan that the processing of autonomic requests by the policy evaluation component in response to blocking conditions in and of themselves can be processed autonomically in the same fashion as noted above in accordance with the rules of a pertinent policy. In this regard, the process of autonomically handling network administration tasks can be viewed as a recursive operation in which interdependent components are managed from the "bottom up". Thus, it will be further apparent to the skilled artisan that the burdensome requirement of a network administrator to remain familiar and aware of all interdependent rules for network administration can be obviated by the autonomic processing of network

administration commands through a set of predefined policy rules which can be processed in accordance with the inventive arrangements.

[0030] The present invention can be realized in hardware, software, or a combination of hardware and software. An implementation of the method and system of the present invention can be realized in a centralized fashion in one computer system, or in a distributed fashion where different elements are spread across several interconnected computer systems. Any kind of computer system, or other apparatus adapted for carrying out the methods described herein, is suited to perform the functions described herein.

[0031] A typical combination of hardware and software could be a general purpose computer system with a computer program that, when being loaded and executed, controls the computer system such that it carries out the methods described herein. The present invention can also be embedded in a computer program product, which comprises all the features enabling the implementation of the methods described herein, and which, when loaded in a computer system is able to carry out these methods.

[0032] Computer program or application in the present context means any expression, in any language, code or notation, of a set of instructions intended to cause a system having an information processing capability to perform a particular function either directly or after either or both of the following a) conversion to another language, code or notation; b) reproduction in a different material form. Significantly, this invention can be embodied in other specific forms without departing from the spirit or essential

attributes thereof, and accordingly, reference should be had to the following claims, rather than to the foregoing specification, as indicating the scope of the invention.